

1995117719

N95-24139

A Terabyte Linear Tape Recorder

John C. Webber
Interferometrics Inc.
8150 Leesburg Pike
Vienna, VA 22182
(703) 790-8500
webber@interf.com

43475
p. 5

A plan has been formulated and selected for a NASA Phase II SBIR award for using the VLBA tape recorder for recording general data. The VLBA tape recorder is a high-speed, high-density linear tape recorder developed for Very Long Baseline Interferometry (VLBI) which is presently capable of recording at rates up to 2 Gbit/sec and holding up to 1 Terabyte of data on one tape, using a special interface and not employing error correction. A general-purpose interface and error correction will be added so that the recorder can be used in other high-speed, high-capacity applications.

The VLBA recorder was developed specifically for recording VLBI data using the Very Long Baseline Array of radio astronomy antennas built by the National Radio Astronomy Observatory. It is an evolution of the technology developed for the NASA Mark IIIA VLBI recording system at MIT Haystack Observatory. Its characteristics may be summarized as follows:

Recording medium:	1-inch-wide D1-equivalent 16 m thick tape
Head type:	38 μ m width, single-crystal ferrite
Bit density:	56,000 flux transitions per inch
Format:	continuous linear tracks, NRZM, magnetic saturation
Tape speed:	For 9 Mbit/sec data, 160 inches per second
MTBF:	10,000 hours typical
Head life:	5,000-20,000 hours depending on environment
Head replacement cost:	\$1/hour
Maintenance:	headstack and tape path cleaned with a cotton swab at each tape change.

A single headstack writes and reads 34 data tracks at a time. The heads are 38 μ m wide and are separated by 698.5 μ m, so that potentially $698.5/38 = 18.38$ passes could be written on the tape. However, some allowance for guard bands between tracks must be made, since the magnetic gap is exactly perpendicular to the direction of tape motion and there must be no crosstalk between tracks. A practical limit is 16 passes, which gives a track spacing of 43.7 μ m with a guard band of about 5 μ m. There are thus 544 data tracks on the tape. Future improvements in tapes and heads are expected to increase this number.

The tape is a D1- or S-VHS equivalent tape available from both 3M and Sony. This tape is 16 m thick, and 20,500 usable feet are contained on a 16-inch reel (only 14-inch reels are currently used). The bit density supportable on this tape is 60,000 bits per inch, so each track contains 14.4 Gigabits. The whole tape with 544 tracks then holds about 8.03 Terabits, or one Terabyte. Using the error-correcting format to be developed, this becomes 788 Gbytes of user data. The cost of one reel of this tape is presently about

\$1500, or \$1.90 per Gigabyte. At the maximum user data rate supported by one headstack, namely 69.44 Mbyte/sec, one tape lasts approximately 192 minutes. Up to 4 headstacks may be mounted on one transport, yielding an aggregate recording rate of 278 Mbyte/sec with a recording time of 48 minutes. This is also the time required to duplicate a tape.

Typical tape life is several hundred read-write cycles including shipping once per month in uncontrolled conditions.

INTERFACES

The data interface in the present VLBA recorder is a set of parallel data lines, each supplying data directly to a single head in the headstack. Formatting of the user data consists of adding synchronization words, time codes, identification data, and parity bits in an external formatter. These bits are simply transferred directly to tape. On playback, the signals from each head are amplified, equalized, and routed to bit synchronizers which recover the clocks from each data stream. All further processing is performed in an external unit which recovers the synchronization codes, de-skews the tracks, and combines the data into a desired format.

In the system under development, the recorder will be responsible for all functions of formatting and deformatting. The user will supply data over a standard interface and recover data from the recorder over the same interface.

For the data rates of concern, there is a prime "standard" interface, namely the High-Performance Parallel Interface, or HIPPI, defined in ANSI X3.183-1991, as defined by the ANSI Task Group X3T9.3. We have adopted this interface as the standard recorder data interface for both record and playback for all applications. The HIPPI channel consists of 32 balanced ECL signals with a common 25 MHz clock and a transfer protocol allowing bursty transmissions. For applications requiring less than the full channel capacity, the recording speed may be varied or the data padded with dummy data to be stripped at playback time.

The HIPPI channel is a one-way device, so two HIPPI channels are needed in order to accommodate both record and playback functions. Commercially available chips provide a single 100 Mbyte/sec HIPPI interface. In the prototype system, two headstacks will be employed, enabling a maximum of 138.88 Mbyte/sec to be recorded. Since the recorder speed can be set with very fine resolution, it will be chosen such that the bit density remains at 60,000 bits per inch.

Since other high-speed protocols and fiber optic interfaces such as advanced ATM are coming into use, it is essential that the recorder be expandable to accommodate them. The plan is to add, for example, an ATM-to-HIPPI interface and continue to operate the recorder exclusively from the HIPPI interface. This simplifies the interfacing problem by placing it outside the recorder proper.

A single low-speed interface will suffice to set the recording mode and control the playback process. This is envisioned to be a 9600 baud RS-232C interface, permitting operation by any computer.

ERROR CORRECTION

For an individual tape track, which is the minimal recording sub-channel, the important characteristics for the VLBA recorder as it is presently used are as follows:

- Random errors: Bit Error Rate (BER) $< 3 \times 10^{-4}$ with 10-year-old tape;
typically 3×10^{-6} with new tape, 3×10^{-5} with 3-year-old tape
- Burst errors: Typical length 100-1000 bits; typical rate 1 burst in 10^7 bits;
bursts usually cause loss of bit synchronization

Random errors are caused by low signal-to-noise ratio (SNR), imperfect equalization, and imperfect bit synchronization. Burst errors are caused primarily by tape defects and are consequently highly dependent on the particular tape in use; any system of error correction must accommodate the worst-case tape. In order to minimize the data lost to dropouts, the distance between sync words should be comparable to the size of the dropouts.

For typical imaging data, a bit error rate better than 1×10^{-9} is required. Other applications require bit error rates as low as 1×10^{-13} .

Recently, VLSI chip sets which implement Reed-Solomon error correction algorithms have become available, and some can run at the data rates in which we are interested. Such chips are available from such manufacturers as LSI Logic and CNR, Inc. It appears that suitable devices for this application are available from Advanced Hardware Architectures (AHA). In particular, the AHA4010 device is attractive because of the following:

Input data format:	8-bit parallel (byte organization)
Data rate:	10 Mbyte/sec (80 Mbit/sec)
Max block length:	255 bytes (programmable)
Max errors correctable:	10 per block (or 20 erasures)
Other features:	No external buffers or control required
Cost:	\$30 each in large quantities

Data will first be coded into data blocks of length 235 bytes. For each such block, we will program the AHA4010 to add 20 error correction bytes to make the total block length 255. This is an overhead of 8.5%.

This scheme will permit correction of up to 10 errors or 20 erasures (or a linear combination of both). Even with a raw bit error rate of 1×10^{-4} , the corrected code block error rate is predicted to be 7×10^{-15} , which satisfies our requirements. The errors will be decorrelated by interleaving the data over a distance long enough so that only one or two bytes from each code block are contained within a single error burst.

CONFIGURATION

The VLBA recorder uses a Metrum Model 96 tape transport, which is a full-size rack of hardware. In the new configuration, it will contain:

- VME-based control computer
- HIPPI interfaces with buffers
- Error correction/formatting boards

- Analog write drivers
- Analog playback amplifiers/equalizers
- Clock recovery (bit sync) boards
- Deskewing buffers
- Transport and headstack motion controllers
- Power supplies

The prototype system will be equipped with two headstacks and so will be capable of recording up to about 138 Mbyte/sec of user data.

RECORDING AND PLAYBACK PROCESS

In addition to error correction, some formatting must be introduced in the form of synchronization words, modulation, and longitudinal parity. Sync words are absolute identifiers of tape block start points. Data modulation by a pseudo-random sequence will guarantee roughly equal numbers of ones and zeroes in the data regardless of content. The addition of longitudinal parity bits on each track will force sufficient transitions in the NRZM format so that good bit synchronizer performance can be maintained.

Upon playback, the parity, modulation, sync, and other formatting information must be undone and stripped out, and the error correction bytes used to restore the original user data. This data must then be reformatted so that it can be transmitted over the HIPPI interface back to the user. A summary of the recording and playback process is shown as Figure 1.

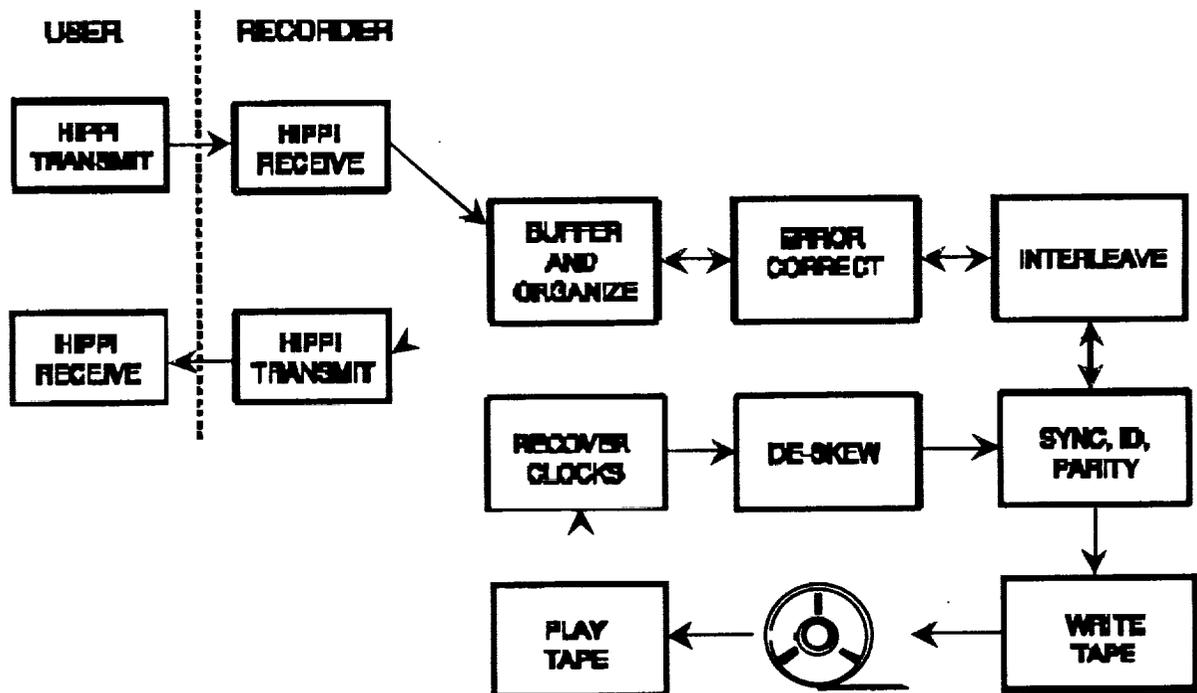


Figure 1. Data flow to and from the recorder

In much of the recording and playback electronics, the same circuitry can be used to perform both the recording and playback processing.

REFERENCES

"A High Data Rate Recorder for Astronomy", H. F. Hinteregger, A. E. E. Rogers, R. J. Cappallo, J. C. Webber, W. T. Petrachenko, and H. Allen, *IEEE Transactions on Magnetics*, **27**, No. 3, p.3455 (1991).

"Very Long Baseline Radio Interferometry: The Mark III System for Geodesy, Astrometry, and Aperture Synthesis", A. E. E. Rogers *et al.*, *Science*, **219**, p.51 (1983).

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE March 1995	3. REPORT TYPE AND DATES COVERED Conference Publication	
4. TITLE AND SUBTITLE Fourth Goddard Conference on Mass Storage Systems and Technologies			5. FUNDING NUMBERS 505	
6. AUTHOR(S) Benjamin Kobler and P. C. Hariharan				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS (ES) Goddard Space Flight Center Greenbelt, Maryland 20771			8. PERFORMING ORGANIZATION REPORT NUMBER 95B00056 Code 505	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS (ES) National Aeronautics and Space Administration Washington, DC 20546-0001			10. SPONSORING / MONITORING AGENCY REPORT NUMBER NASA CP-3295	
11. SUPPLEMENTARY NOTES Kobler: Goddard Space Flight Center, Greenbelt, MD; Hariharan: Systems Engineering and Security, Inc., Greenbelt, Maryland				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Unclassified - Unlimited Subject Category 82 Availability: NASA CASI (301) 621-0390.			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) This report contains copies of all those technical papers received in time for publication just prior to the Fourth Goddard Conference on Mass Storage Systems and Technologies, held March 28-30, 1995, at the University of Maryland, University College Conference Center, in College Park, Maryland. This series of conferences continues to serve as a unique medium for the exchange of information on topics relating to the ingestion and management of substantial amounts of data and the attendant problems involved. This year's discussion topics include new storage technology, stability of recorded media, performance studies, storage system solutions, the National Information infrastructure (Infobahn), the future for storage technology, and lessons learned from various projects. There also will be an update on the IEEE Mass Storage Systems Reference Model Version 5, on which the final vote was taken in July 1994.				
14. SUBJECT TERMS Magnetic tape, magnetic disk, optical disk, mass storage, software storage, digital recording, data compression			15. NUMBER OF PAGES 408	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL	

